arXiv:1704.03864v1 [math.PR] 12 Apr 2017

Matrix Concentration for Expander Walks

Ankit Garg* Nikhil Srivastava†

April 13, 2017

Abstract

We prove a Chernoff-type bound for sums of matrix-valued random variables sampled via a random walk on an expander, confirming a conjecture of Wigderson and Xiao up to logarithmic factors in the deviation parameter. Our proof is based on a recent multi-matrix extension of the Golden-Thompson inequality due to Sutter et al. [SBT17], as well as an adaptation of an argument for the scalar case due to Healy [Hea08]. Secondarily, we also provide a generic reduction showing that any concentration inequality for vector-valued martingales implies a concentration inequality for the corresponding expander walk, with a weakening of parameters proportional to the squared mixing time.

1 Introduction

The Chernoff Bound [Che52] is one of the most widely used probabilistic results in computer science. It states that a sum of independent bounded random variables exhibits subgaussian concentration around its mean. In particular, when the random variables are i.i.d. samples from a fixed distribution, it implies that the empirical mean of k samples is ϵ-close to the true mean with exponentially decaying probability proportional to $e^{-\Omega(k\epsilon^2)}$.

An important generalization of this bound was achieved by Gillman [Gil98] (with refinements later by [Lez98, Kah97, LP04, WX05, Hea08, Wag08, CLLM12, RR17]), who significantly relaxed the independence assumption. In particular, suppose $X : V \to \mathbb{C}$ is a bounded random variable supported on a finite probability space V, G is a regular expander graph with vertex set V of size n, and $v_1, \ldots, v_k$ is a stationary random walk of length k on G. Then, even though the random variables $X(v_i)$ are in not independent (unless G is the complete graph), it is shown that:

$$\mathbb{P}\left[\left|\frac{1}{k}\sum_{i=1}^{k} X(v_i) - \mathbb{E}X\right| > \epsilon\right] \leq 2 \cdot \exp(-\Omega((1-\lambda)k\epsilon^2)), \tag{1}$$

where $1-\lambda$ is the spectral gap of the transition matrix of the random walk. The gain here is that sampling a stationary random walk of length k on a constant degree graph requires $\log(n) + O(k)$ random bits, which is much less than the $k\log(n)$ bits required to produce k independent samples. Since constant degree expanders with constant spectral gap can be explicitly constructed

*Microsoft Research New England, garga@microsoft.com

†UC Berkeley, nikhil@math.berkeley.edu. This work was partially supported by NSF grant CCF-1553751 and a Sloan Research Fellowship.

[RVW00, LPS88], this leads to a generic "derandomization" of the Chernoff bound, which has had several important applications (see the introduction of [WX05] for a discussion).

A different generalization of the Chernoff bound appeared in the works of Rudelson [Rud99], Ahlswede-Winter [AW02], and Tropp [Tro12], who showed that a similar concentration phenomenon is true for *matrix-valued* random variables. In particular, if $X_1, \ldots, X_k$ are independent $d \times d$ complex Hermitian random matrices with $\|X_i\| \leq 1$, then the following is true:

$$\mathbb{P}\left[\left\|\frac{1}{k}\sum_{i=1}^{k} X_i - \mathbb{E}X\right\| > \epsilon\right] \leq 2d \cdot \exp(-\Omega(k\epsilon^2)). \tag{2}$$

The only difference between this and the usual Chernoff bound is the factor of d in front of the deviation probability; to see that it is necessary, notice that the diagonal case simply corresponds to a direct sum of d arbitrarily correlated instances of the scalar Chernoff bound, so by the union bound the probability should be d times as large in the worst case. This so called "Matrix Chernoff Bound" has seen several applications as well, notably in quantum information theory, numerical linear algebra, and spectral graph theory; the reader may consult e.g. the book [Tro15] for many examples.

We present two different extensions of the above results in this paper.

1.1 A Matrix Expander Chernoff Bound

It is natural to wonder whether there is a common generalization of (1) and (2), i.e., a "Matrix Expander Chernoff Bound". Such a result was conjectured by Wigderson and Xiao in [WX] — in fact, [WX05] contained a proof of it, but the authors later discovered a gap in the proof. In this paper, we prove Wigderson and Xiao's conjecture up to a factor of $\log^2(1/\epsilon)$, namely:

Theorem 1.1. *Let G be a regular λ-expander on V and let f be a function $f : V \to M_d(\mathbb{C})$ s.t.*

1. *For each v, $f(v)$ is Hermitian and $\|f(v)\| \leq 1$.*
2. *$\sum_{v \in V} f(v) = 0$.*

Then, for a stationary random walk $v_1, \ldots, v_k$ with $k \geq \log(d)$ and $\epsilon \in (0,1)$ we have:

$$\mathbb{P}\left[\lambda_{\max}\left(\sum_{j=1}^{k} f(v_j)\right) \geq k\epsilon\right] \leq d \cdot \exp\left(-\Omega\left(\frac{\epsilon^2(1-\lambda)k}{\log^2(1/\epsilon(1-\lambda))}\right)\right)$$

$$\mathbb{P}\left[\lambda_{\min}\left(\sum_{i=1}^{k} f(v_j)\right) \leq -k\epsilon\right] \leq d \cdot \exp\left(-\Omega\left(\frac{\epsilon^2(1-\lambda)k}{\log^2(1/\epsilon(1-\lambda))}\right)\right)$$

The dependence on ϵ here is off by a factor of roughly $\log^2(1/\epsilon)$ from the dependence in Chernoff-type bounds. However, in most complexity-theoretic applications, ϵ is taken to be a constant, so this additional dependence is not very consequential. This theorem adds to the amazingly long list of pseudorandom properties of expander graphs. We suspect that it will find further applications in the study of derandomization and expander graphs, beyond the ones mentioned in [WX05] (which were later recovered by different methods in [WX08]). We believe the logarithmic factor can be removed but do not know how to do that at the moment, and leave it as an open problem.

To describe the ideas that go into the proof, let us begin by recalling how the usual scalar Chernoff bound is proved, in the case when the random variables have mean zero. The key observation is that if $X_1, \ldots, X_k$ are independent random variables, then the moment generating function of the sum is equal to the product of the moment generating functions:

$$\mathbb{E} \exp\left(t \sum_{i=1}^{k} X_i\right) = \prod_{i=1}^{k} \mathbb{E} \exp(tX_i).$$

This is no longer true in case where the X_i come from a random walk, but we still have the algebraic fact that

$$\exp\left(t \sum_{i=1}^{k} X_i\right) = \prod_{i=1}^{k} \exp(tX_i), \tag{3}$$

which allows one to decompose the sum as a product. The latter allows one to consider the steps of the random walk separately and analyze the change in the expectation inductively.

The analogue of the moment generating function in the matrix case is

$$\mathbb{E}\mathrm{tr}\left[\exp\left(t \sum_{i=1}^{k} X_i\right)\right],$$

and the main difficulty is that (3) no longer holds if the matrices X_i do not commute. A substitute for this fact is given by the Golden-Thompson inequality [Gol65, Tho65], which states that for any Hermitian A, B:

$$\mathrm{tr}(\exp(A+B)) \leq \mathrm{tr}(\exp(A)\exp(B)). \tag{4}$$

The latter expression may further be bounded by $\|\exp(A)\|\mathrm{tr}(\exp(B))$, and this is sufficient to prove (2) in the independent case as is done in [AW02], where an inductive application of it yields

$$\mathbb{E}\mathrm{tr}\left[\exp\left(t \sum_{i=1}^{k} X_i\right)\right] \leq \mathrm{tr}(I) \cdot \prod_{i=1}^{k} \left\|\mathbb{E} \exp(tX_i)\right\|.$$

However, this approach is too crude to handle the Markov case, roughly because in the absence of inependence, passing to the norm makes it difficult to utilize the fact that the expectation of each X_i is zero.

The original proof of Wigderson-Xiao was based on the following plausible multimatrix generalization of (4):

$$\mathrm{tr}\left[\exp\left(\sum_{i=1}^{k} A_i\right)\right] \leq \mathrm{tr}\left[\prod_{i=1}^{k} \exp(A_i)\right],$$

which turns out to be false for $k > 2$. To see why, observe that the left hand side is always nonnegative, whereas the right hand side can be the trace of a product of any three positive semidefinite matrices, which can be negative (and this is not the case for two matrices). This led to a fatal gap in their proof.

The main technical ingredient in our proof is the following recent multimatrix generalization of (4) due to [SBT17] (see also [HKT16]):

Theorem 1.2 (Corollary 3.3 in [SBT17]). *Let $H_1, \ldots, H_k \in M_d(\mathbb{C})$ be Hermitian matrices. Then*

$$\log\left[\operatorname{tr}\left(\exp\left(\sum_{j=1}^{k} H_j\right)\right)\right] \leq \int_{-\infty}^{\infty} \beta(b) \log\left[\operatorname{tr}\left(\prod_{j=1}^{k} \exp\left(\frac{H_j(1+ib)}{2}\right) \prod_{j=k}^{1} \exp\left(\frac{H_j(1-ib)}{2}\right)\right)\right] \mathrm{d}b$$

Here β is a distribution given by $\beta(b) = \frac{\pi}{\exp(\pi b)+\exp(-\pi b)+2}$.

The notable features of this theorem are:

(a) Since $\exp(H(1+ib)) = \exp(H(1-ib))^\dagger$ for Hermitian H, The right hand always considers the trace of a matrix times its adjoint, which is always positive semidefinite, ruling out the bad example described above.

(b) It is an *average case* inequality, in the sense that it holds in an averaged sense after multiplying the $\exp(H_j)$ by certain random unitaries $\exp(ibH_j)$, where the b comes from a specific universal distribution. We remark that the inequality is known to be false in the worst case (i.e., with $b = 0$ and with other small values of b; see [SBT17] for a discussion).

(c) The average is taken over the logarithm of the trace, which turns out to be crucial in our proof.

Theorem 1.2 allows one to relate the exponential of a sum of matrices to a (two-sided) product of $d \times d$ matrices and their adjoints. In order to mimic the strategy of Wigderson-Xiao, we rewrite this as a one-sided matrix product of $d^2 \times d^2$ matrices acting on $\mathbb{C}^{d \times d}$, by encoding left and right multiplication on this space via a tensor product. However, these matrices are no longer Hermitian or bounded (due to the fact that the distribution β has unbounded support). We surmount these difficulties by using a variant of the more robust linear algebraic proof technique of Healy [Hea08] (as opposed to the perturbation theoretic approach of [WX05], which is difficult to apply here because the relevant matrices are not normal), and showing that the distribution of β can be truncated without much loss; this truncation is the source of the extra $\log^2(1/\epsilon)$ term which appears in our final bound. The proof of Theorem 1.1 is presented in Section 3.

Remark 1.3. The peculiar form of the inequality 1.2 arises from the fact that its proof relies on Hirschman's interpolation theorem [Hir52] in complex analysis, which ultimately boils down to the fact that the value of a holomorphic function at any point inside a domain is a suitable average of the values that it takes on the boundary. In the present context that distribution turns out to be (a transformation of) the Poisson kernel, which becomes the distribution β seen above. The interested reader is encouraged to consult [SBT17] for a more detailed discussion.

1.2 Martingale Approximation of Expander Walks

While Theorem 1.1 provides a satisfactory generalization of the expander Chernoff bound to the case when one is interested in the spectral norm of a matrix-valued function on V, one could ask what happens for other matrix norms (such as Schatten norms), or even more generally, for functions taking values in an arbitrary Banach space. Our second contribution is a generic reduction from this problem, of proving concentration for random variables sampled using a *Markov chain*, to the much more well-studied problem (see e.g. [CL06]) of concentration for sums of *martingale* random variables.

Theorem 1.4. *Suppose $G = (V, E)$ is a λ-expander and $f : V \to \mathbb{R}^N$ is a vector-valued function satisfying $\sum_{v \in V} f(v) = 0$ with $F := \sqrt{\sum_{v \in V} \|f(v)\|_2^2}$. If $v_1, \ldots, v_k$ is a stationary random walk on G, then for every $\epsilon > 0$, there is a martingale difference sequence $Z_1, \ldots, Z_k$ with respect to the filtration generated by initial segments of $v_1, \ldots, v_k$ such that*

$$\frac{1}{k} \sum_{i=1}^{k} f(v_i) = W + \frac{1}{k} \sum_{i=1}^{k} Z_i,$$

where

1. *W is a random vector satisfying $\|W\|_2 \leq \epsilon$.*
2. *Each term Z_i satisfies*

$$\|Z_i\|_* \leq \frac{2 \log(F/\epsilon)}{1 - \lambda} \cdot \max_{v \in V} \|f(v)\|_*,$$

for every norm $\|\cdot\|_*$.

Thus, the empirical sums of any bounded (in any norm) function on a graph are well-approximated by a martingale whose increments are also bounded, with a loss in the bound depending on the ℓ_2 norm F of the function and the spectral gap of the graph. Since F will typically scale with the number of vertices, the ratio above is typically comparable to the mixing time.

To see the theorem in action, consider the case when $f(v)$ is matrix-valued in $d \times d$ Hermitian matrices and $\|\cdot\|_*$ is the operator norm. If $\|f(v)\|_{op} \leq 1$ then we have the bound

$$F^2 = \sum_{v \in V} \|f(v)\|_F^2 \leq dn.$$

Suppose we are interested in obtaining an estimate on the probability:

$$\mathbb{P}\left[\left\|\frac{1}{k} \sum_{v \in V} f(v)\right\|_{op} > \epsilon\right]. \tag{5}$$

Applying Theorem 1.4 with parameter $\epsilon/2$ and noting that $\|W\|_{op} \leq \|W\|_F$, we have that (5) is at most

$$\mathbb{P}\left[\left\|\frac{1}{k} \sum_{i=1}^{k} Z_i\right\| > \epsilon/2\right],$$

where Z_i is a martingale with bound

$$\|Z_i\|_{op} \leq \frac{\log(nd/\epsilon)}{1 - \lambda}.$$

We now appeal to the existing martingale generalization of (2) (see e.g. [Tro12]) and find that this probability is at most

$$2d \cdot \exp\left(-\Omega\left(\frac{k\epsilon^2(1-\lambda)^2}{\log^2(nd)}\right)\right).$$

While this theorem is much weaker than the previous one in terms of parameters (depending on the square of the mixing time rather than on the spectral gap), it shows qualitatively that

concentration for Markov chains is a generic phenomenon rather than something specific to matrices. It also allows one to instantly import the wealth of results regarding concentration for martingales in various Banach spaces (see e.g., [LT13]) to the random walk setting, albeit with suboptimal parameters. The simple proof of Theorem 1.4 is presented in Section 4.

1.3 Acknowledgments

We thank Yash Deshpande, Ravi Kannan, Zeph Landau, Avi Wigderson and David Xiao for helpful conversations, as well Microsoft Research India and the Simons Institute for the Theory of Computing, where much of this work was carried out.

2 Preliminaries

We will be working with D-regular undirected graphs $G = (V, E)$. The number of vertices of the graph, V will be denoted by n. "A" will denote the adjacency matrix of the graph and $P = A/D$ will denote its normalized adjacency matrix. A regular graph G will be called a λ-expander $(0 < \lambda < 1)$ if $\|P\mathbf{x}\| \leq \lambda \cdot \|\mathbf{x}\|$ for all vectors $\mathbf{x} \in \mathbb{C}^n$ s.t. $\sum_{i=1}^{n} \mathbf{x}_i = 0$.

We will use $e_i \in \mathbb{C}^d$ to denote the standard basis vector with 1 in i^{th} position and 0 everywhere else. $M_d(\mathbb{C})$ will denote the set of $d \times d$ complex matrices.

Given two vectors $v \in \mathbb{C}^{d_1}$ and $w \in \mathbb{C}^{d_2}$, their tensor product $v \otimes w \in \mathbb{C}^{d_1 d_2}$ is the vector whose $(i, j)^{\text{th}}$ entry is $v(i)w(j)$ (for concreteness, assume the entries are in lexicographic order). Given two matrices $A_1 \in M_{d_1}(\mathbb{C})$ and $A_2 \in M_{d_2}(\mathbb{C})$, their tensor product $A_1 \otimes A_2 \in M_{d_1 d_2}(\mathbb{C})$ is the matrix whose $((i, k), (j, l))^{\text{th}}$ entry is $A_1(i, j)A_2(k, l)$. It is easy to see that

$$(A \otimes B)v \otimes w = Av \otimes Bw$$

and

$$(A \otimes B)(C \otimes D) = AB \otimes CD.$$

For a matrix $X \in M_d(\mathbb{C})$, $\text{vec}(X) \in \mathbb{C}^{d^2}$ will denote the vectorized version of the matrix X. That is

$$\text{vec}(X) = \sum_{i,j=1}^{d} X(i, j)e_i \otimes e_j$$

We have the following relationship between matrix multiplication and the tensor product:

$$\text{vec}(AXB) = (A \otimes B^T)\text{vec}(X).$$

All logarithms will be taken with the base e, and $\exp(x)$ will denote e^x. The matrix exponential of a complex matrix $X \in M_d(\mathbb{C})$ is defined by the Taylor expansion:

$$\exp(X) = \sum_{j=0}^{\infty} \frac{X^j}{j!}.$$

The series converges for all matrices X. We will use the fact that

$$\exp(A) \otimes \exp(B) = \exp(A \otimes I + I \otimes B),$$

which may be checked by expanding both sides and comparing terms.

The following simple proposition will be useful in our proofs.

Proposition 2.1. *Let A and B be Hermitian psd matrices. Then*

$$\mathrm{tr}[AB] \le \|A\| \cdot \mathrm{tr}[B]$$

Proof. Let $B = \sum_{j=1}^{n} \sigma_j v_j v_j^\dagger$ be the eigenvalue decomposition of B. Then

$$\mathrm{tr}[AB] = \sum_{j=1}^{n} \sigma_j \mathrm{tr}\left[A v_j v_j^\dagger\right] = \sum_{j=1}^{n} \sigma_j v_j^\dagger A v_j \le \sum_{j=1}^{n} \sigma_j \cdot \|A\| = \|A\| \cdot \mathrm{tr}[B]$$

□

3 Proof of Theorem 1.1

In this section we present the proof of Theorem 1.1. We restate it here (with explicit constants) for convenience.

Theorem 3.1. *Let G be a regular λ-expander on V and let f be a function $f : V \to M_d(\mathbb{C})$ s.t.*

1. *For each v, $f(v)$ is Hermitian and $\|f(v)\| \le 1$.*
2. *$\sum_{v \in V} f(v) = 0$.*

If $v_1, \ldots, v_k$ is a stationary random walk on G, then for $k \ge \log(d)$ and $\epsilon \in (0,1)$,

$$\mathbb{P}\left[\lambda_{\max}\left(\sum_{j=1}^{k} f(v_j)\right) \ge k\epsilon\right] \le d \cdot \exp\left(-\frac{\epsilon^2(1-\lambda)k}{1200\log^2(150/\epsilon(1-\lambda))}\right)$$
$$\mathbb{P}\left[\lambda_{\min}\left(\sum_{i=1}^{k} f(v_j)\right) \le -k\epsilon\right] \le d \cdot \exp\left(-\frac{\epsilon^2(1-\lambda)k}{1200\log^2(150/\epsilon(1-\lambda))}\right)$$

Proof. Due to symmetry, it suffices to prove just one of the statements. Let $t > 0$ be a parameter to be chosen later. Then

$$\begin{aligned}\mathbb{P}\left[\lambda_{\max}\left(\sum_{j=1}^{k} f(v_j)\right) \ge k\epsilon\right] &\le \mathbb{P}\left[\mathrm{tr}\left(\exp\left(t\sum_{j=1}^{k} f(v_j)\right)\right) \ge \exp(tk\epsilon)\right] \\ &\le \frac{\mathbb{E}\left[\mathrm{tr}\left(\exp\left(t\sum_{j=1}^{k} f(v_j)\right)\right)\right]}{\exp(tk\epsilon)} \end{aligned} \tag{6}$$

The second inequality follows from Markov's inequality. Now by Theorem 1.2, for every fixed walk $v_1, \ldots, v_k$, we have the deterministic inequality:

$$\log\left[\mathrm{tr}\left(\exp\left(t\sum_{j=1}^{k} f(v_j)\right)\right)\right] \le \int_{-\infty}^{\infty} \beta(b) \log\left[\mathrm{tr}\left(\prod_{j=1}^{k} \exp\left(\frac{tf(v_j)(1+ib)}{2}\right)\prod_{j=k}^{1} \exp\left(\frac{tf(v_j)(1-ib)}{2}\right)\right)\right] \mathrm{d}b \tag{7}$$

Also note that for all $v_1, \ldots, v_k$:

$$\begin{aligned}
&\mathrm{tr}\left(\prod_{j=1}^{k}\exp\left(\frac{tf(v_j)(1+ib)}{2}\right)\prod_{j=k}^{1}\exp\left(\frac{tf(v_j)(1-ib)}{2}\right)\right)\\
&\leq \|\exp(t(f(v_1))\|\cdot\mathrm{tr}\left(\prod_{j=2}^{k}\exp\left(\frac{tf(v_j)(1+ib)}{2}\right)\prod_{j=k}^{2}\exp\left(\frac{tf(v_j)(1-ib)}{2}\right)\right)\\
&\ldots\\
&\leq d\cdot\exp(tk) \qquad (8)
\end{aligned}$$

by inductively applying cyclicity of the trace and Proposition 2.1, and noting that $\|\exp(tf(v))\| \leq \exp(t)$ for all $v \in V$. Thus, for any $\alpha > 0$ (a parameter to be set later), we can bound

$$\begin{aligned}
&\int_{-\infty}^{\infty}\beta(b)\log\left[\mathrm{tr}\left(\prod_{j=1}^{k}\exp\left(\frac{tf(v_j)(1+ib)}{2}\right)\prod_{j=k}^{1}\exp\left(\frac{tf(v_j)(1-ib)}{2}\right)\right)\right]\mathrm{d}b\\
&\leq\int_{-\alpha}^{\alpha}\beta(b)\log\left[\mathrm{tr}\left(\prod_{j=1}^{k}\exp\left(\frac{tf(v_j)(1+ib)}{2}\right)\prod_{j=k}^{1}\exp\left(\frac{tf(v_j)(1-ib)}{2}\right)\right)\right]\mathrm{d}b+\int_{|b|\geq\alpha}\beta(b)(kt+\log(d))\mathrm{d}b\\
&\leq\log\left[\int_{-\alpha}^{\alpha}\tilde{\beta}(b)\mathrm{tr}\left(\prod_{j=1}^{k}\exp\left(\frac{tf(v_j)(1+ib)}{2}\right)\prod_{j=k}^{1}\exp\left(\frac{tf(v_j)(1-ib)}{2}\right)\right)\mathrm{d}b\right]+2(kt+\log(d))\exp(-\pi\alpha)
\end{aligned} \qquad (9)$$

The first inequality follows from equation (8), and the second inequality follows from the concavity of log and the fact that $\beta(b) \leq \min\left\{\frac{\pi}{\exp(\pi b)}, \frac{\pi}{\exp(-\pi b)}\right\}$. Here $\tilde{\beta}$ is the distribution β restricted to the interval $[-\alpha, \alpha]$, and this truncation will be crucial later on in the proof.

Choose $\alpha = 1 + \log(1/t)$ to ensure that $2(kt + \log(d))\exp(-\pi\alpha) \leq 2kt^2$ (this uses the fact that $k \geq \log(d)$). Then combining equations (7) and (9), we get that for every walk $v_1, \ldots, v_k$:

$$\left[\mathrm{tr}\left(\exp\left(t\sum_{j=1}^{k}f(v_j)\right)\right)\right]\leq\exp(2kt^2)\cdot\int_{-\alpha}^{\alpha}\tilde{\beta}(b)\mathrm{tr}\left(\prod_{j=1}^{k}\exp\left(\frac{tf(v_j)(1+ib)}{2}\right)\prod_{j=k}^{1}\exp\left(\frac{tf(v_j)(1-ib)}{2}\right)\right)\mathrm{d}b \qquad (10)$$

We now take the expectation over all walks to obtain the following upper bound on the moment generating function in (6):

$$\mathbb{E}\left[\mathrm{tr}\left(\exp\left(t\sum_{j=1}^{k}f(v_j)\right)\right)\right]\leq\exp(2kt^2)\cdot\max_{|b|\leq\alpha}\mathbb{E}\mathrm{tr}\left(\prod_{j=1}^{k}\exp\left(\frac{tf(v_j)(1+ib)}{2}\right)\prod_{j=k}^{1}\exp\left(\frac{tf(v_j)(1-ib)}{2}\right)\right).$$

The core of the proof is the following bound on the moment generating function-like expression that appears above, for sufficiently small b:

Lemma 3.2. *Suppose $0 < t \leq \min\{1/4, \log(1/\lambda)/2\}$. Then for $|b| \leq 1 + \log(1/t)$,*

$$\mathbb{E}\mathrm{tr}\left(\prod_{j=1}^{k}\exp\left(\frac{tf(v_j)(1+ib)}{2}\right)\prod_{j=k}^{1}\exp\left(\frac{tf(v_j)(1-ib)}{2}\right)\right)\leq d\cdot\exp\left(k\left(9+\frac{72\lambda}{1-\lambda}\right)t^2\log^2(1/t)\right)$$

Assuming this lemma, we can easily complete the proof of the theorem as:

$$\mathbb{E}\left[\mathrm{tr}\left(\exp\left(t\sum_{j=1}^{k} f(v_j)\right)\right)\right] \leq d\cdot\exp(2kt^2)\cdot\exp\left(k\left(9+\frac{72\lambda}{1-\lambda}\right)t^2\log^2(1/t)\right)$$
$$\leq d\cdot\exp\left(k\left(\frac{72}{1-\lambda}\right)t^2\log^2(1/t)\right)$$

Combining this with equation (6), we get that

$$\mathbb{P}\left[\lambda_{\max}\left(\sum_{j=1}^{k} f(v_j)\right)\geq k\epsilon\right] \leq d\cdot\exp\left(\left(\frac{72}{1-\lambda}\right)kt^2\log^2(1/t)-kt\epsilon\right)$$

Choose t s.t. $\left(\frac{72}{1-\lambda}\right)t^2\log^2(1/t)=t\epsilon/2$. This ensures that $t\leq\min\{1/4,\log(1/\lambda)/2\}$ and also that

$$\mathbb{P}\left[\lambda_{\max}\left(\sum_{j=1}^{k} f(v_j)\right)\geq k\epsilon\right] \leq d\cdot\exp(-tk\epsilon/2)$$

Now $t\log^2(1/t)=\delta$ implies that $t\geq\frac{\delta}{4\log^2(1/\delta)}$ for a small enough constant δ. This gives the required bound. □

We now give the proof of Lemma 3.2

Proof of Lemma 3.2. We start by writing the expected trace expression in terms of the transition matrix of the random walk. This is an analogue of a step which is common to most of the expander chernoff bound proofs in the scalar case. Let P be the normalized adjacency matrix of G and let $\widetilde{P}=P\otimes I_{d^2}$. Let E denote the $nd^2\times nd^2$ block diagonal matrix where the v^{th} diagonal block is the matrix

$$M_v=\exp\left(\frac{tf(v)(1+ib)}{2}\right)\otimes\exp\left(\frac{tf(v)(1-ib)}{2}\right)$$

Then $\left(E\widetilde{P}\right)^k$ is an $nd^2\times nd^2$ block matrix whose $(u,v)^{\text{th}}$ $(d^2\times d^2)$ block is given by the matrix

$$\sum_{v_1,\ldots,v_{k-1}} P_{u,v_1}\cdot\left(\prod_{j=1}^{k-2}P_{v_j,v_{j+1}}\right)\cdot P_{v_{k-1},v}\cdot M_u\cdot\left(\prod_{j=1}^{k-1}M_{v_j}\right)$$

Let $z_0\in\mathbb{C}^{nd^2}$ be the vector $\frac{\mathbf{1}}{\sqrt{n}}\otimes\mathrm{vec}(I_d)$. Here $\mathbf{1}$ is the all 1's vector and $\mathrm{vec}(I_d)$ is the vector form of the identity matrix. Then, by applying

$$\langle\mathrm{vec}(I_d),A_1\otimes A_2\,\mathrm{vec}(I_d)\rangle=\mathrm{tr}\left[A_1A_2^T\right]$$

it follows that for a stationary random walk $v_1,\ldots,v_k$:

$$\mathbb{E}\mathrm{tr}\left(\prod_{j=1}^{k}\exp\left(\frac{tf(v_j)(1+ib)}{2}\right)\prod_{j=k}^{1}\exp\left(\frac{tf(v_j)(1-ib)}{2}\right)\right)=\mathbb{E}\left\langle\mathrm{vec}(I_d),\prod_{i=1}^{k}M_{v_i}\mathrm{vec}(I_d)\right\rangle$$
$$=\left\langle z_0,\left(E\widetilde{P}\right)^k z_0\right\rangle$$

Hence we can focus our attention on $\left\langle z_0, \left(E\widetilde{P}\right)^k z_0 \right\rangle$.

Let $\mathbf{1}$ be the all 1's vector. For a vector $z \in \mathbb{C}^{nd^2}$, let $z^{\|}$ denote the component of z which lies in the subspace spanned by the d^2 vectors $\mathbf{1} \otimes e_i$, $1 \le i \le d^2$. Let $z^{\perp}$ denote the component in the orthogonal space. Letting $z_j = \left(E\widetilde{P}\right)^j z_0$, we are interested in bounding

$$\langle z_0, z_k \rangle = \langle z_0, z_k^{\|} \rangle \le \|z_0\| \cdot \|z_k^{\|}\| = \sqrt{d} \cdot \|z_k^{\|}\|$$

The following lemma is the analogue of the main lemma in Healy's proof for the scalar valued expander chernoff bound [Hea08]. Roughly speaking, it tracks how much a vector can move in and out of the subspace we are interested in as the operator $E\widetilde{P}$ is applied.

Lemma 3.3. *Let G be a regular λ-expander on V. Suppose each vertex v is assigned a matrix $H_v \in M_{d^2}(\mathbb{C})$ s.t. $\|H_v\| \le \ell$ and $\sum_v H_v = 0$. Let P be the normalized adjacency matrix of G and let $\widetilde{P} = P \otimes I_{d^2}$. Let E denote the $nd^2 \times nd^2$ block diagonal matrix where the v^{th} diagonal block is the matrix $\exp(tH_v)$. Also suppose that $\|E\| = \max_v \|\exp(tH_v)\| \le \exp(t)$ and that $0 < t \le \log(1/\lambda)/2$. Then for any $z \in \mathbb{C}^{nd^2}$*

1. $\|(E\widetilde{P}z^{\|})^{\|}\| \le \alpha_1 \|z^{\|}\|$
2. $\|(E\widetilde{P}z^{\|})^{\perp}\| \le \alpha_2 \|z^{\|}\|$
3. $\|(E\widetilde{P}z^{\perp})^{\|}\| \le \alpha_3 \|z^{\perp}\|$
4. $\|(E\widetilde{P}z^{\perp})^{\perp}\| \le \alpha_4 \|z^{\perp}\|$

where

1. $\alpha_1 = (\exp(t\ell) - t\ell)$
2. $\alpha_2 = \exp(t\ell) - 1$
3. $\alpha_3 = \lambda \cdot (\exp(t\ell) - 1)$
4. $\alpha_4 = \sqrt{\lambda}$

Proof of Lemma 3.3. 1. $(E\widetilde{P}z^{\|})^{\|} = (Ez^{\|})^{\|}$. Suppose $z^{\|} = \mathbf{1} \otimes \mathbf{w}$ for some $\mathbf{w} \in \mathbb{C}^{d^2}$. Then $\|z^{\|}\| = \sqrt{n} \cdot \|\mathbf{w}\|$ and

$$(Ez^{\|})^{\|} = \mathbf{1} \otimes \left(\frac{1}{n} \sum_{v \in V} \exp(H_v)\mathbf{w} \right)$$

Thus we are looking to bound $\|\frac{1}{n}\sum_{v\in V} \exp(H_v)\|$. By Taylor's expansion,

$$\begin{aligned} \left\| \frac{1}{n} \sum_{v \in V} \exp(H_v) \right\| &= \left\| I + \sum_{j \ge 2} \frac{1}{n} \sum_{v \in V} \frac{t^j}{j!} H_v^j \right\| \\ &\le 1 + \sum_{j \ge 2} \frac{t^j}{j!} \ell^{j/2} \\ &= \exp(t\ell) - t\ell \end{aligned}$$

The first equality uses the fact that $\sum_v H_v = 0$. The inequality follows from triangle inequality and the fact that $\|H_v\| \le \ell$.

2. Note that
$$(E\widetilde{P}z^{\|})^{\perp} = (Ez^{\|})^{\perp} = ((E-I)z^{\|})^{\perp}$$
since $(z^{\|})^{\perp} = 0$. Hence we can bound
$$\begin{aligned}\left\|((E-I)z^{\|})^{\perp}\right\| \le \left\|(E-I)z^{\|}\right\| &\le \|E-I\| \cdot \|z^{\|}\| \\ &= \max_v \|\exp(tH_v) - I\| \cdot \|z^{\|}\| \\ &= \max_v \left\|\sum_{j\ge 1} \frac{t^j}{j!} H_v^j\right\| \cdot \|z^{\|}\| \\ &\le \left(\sum_{j\ge 1} \frac{t^j}{j!}\ell^j\right) \cdot \|z^{\|}\| \\ &= (\exp(t\ell) - 1) \cdot \|z^{\|}\|\end{aligned}$$
The third inequality follows from triangle inequality and the fact that $\|H_v\| \le \ell$.

3. Note that
$$(E\widetilde{P}z^{\perp})^{\|} = ((E-I)\widetilde{P}z^{\perp})^{\|}$$
This is because $(\widetilde{P}z^{\perp})^{\|} = 0$ since $\widetilde{P}$ preserves the property of being orthogonal to the space spanned by the vectors $\mathbf{1} \otimes e_i$ (these are the top eigenvectors of $\widetilde{P}$). Hence we can bound
$$\begin{aligned}\left\|((E-I)\widetilde{P}z^{\perp})^{\|}\right\| &\le \left\|(E-I)\widetilde{P}z^{\perp}\right\| \\ &\le \|E-I\| \cdot \left\|\widetilde{P}z^{\perp}\right\| \\ &\le (\exp(t\ell) - 1) \cdot \lambda \cdot \left\|z^{\perp}\right\|\end{aligned}$$
Third inequality follows from the fact that $\|E-I\| \le \exp(t\ell) - 1$ and that G is a λ-expander.

4. We can bound
$$\|(E\widetilde{P}z^{\perp})^{\perp}\| \le \|E\widetilde{P}z^{\perp}\| \le \exp(t) \cdot \lambda \cdot \|z^{\perp}\| \le \sqrt{\lambda} \cdot \|z^{\perp}\|$$
Second inequality follows from the fact that $\|E\| \le \exp(t)$ and that G is a λ-expander. Third inequality follows from the fact that $t \le \log(1/\lambda)/2$.

□

We now use Lemma 3.3 to analyze the evolution of $z_j^{\|}$ and $z_j^{\perp}$. Note that we satisfy its conditions with
$$H_v = \frac{f(v)(1+ib)}{2} \otimes I_d + I_d \otimes \frac{f(v)(1-ib)}{2}.$$
In particular, since $\|H_v\| \le \sqrt{1+b^2} \le 1 + |b| \le 2 + \log(1/t)$ we can set $\ell = 2 + \log(1/t)$. Also
$$\|\exp(tH_v)\| = \|\exp(t\,\mathrm{Re}(H_v))\| = \left\|\exp\left(t\left(\frac{f(v)}{2} \otimes I_d + I_d \otimes \frac{f(v)}{2}\right)\right)\right\| \le \exp(t),$$

and $\sum_v H_v = 0$ since $\sum_v f(v) = 0$.

Let us first prove that $\|z_i^\perp\| \le \alpha_5 \cdot \max_{j<i} \|z_j^\parallel\|$, where $\alpha_5 = \frac{\alpha_2}{1-\alpha_4}$.

$$\begin{aligned}
\|z_i^\perp\| &= \left\|(E\widetilde{P} z_{i-1})^\perp\right\| \\
&\le \left\|(E\widetilde{P} z_{i-1}^\parallel)^\perp\right\| + \left\|(E\widetilde{P} z_{i-1}^\perp)^\perp\right\| \\
&\le \alpha_2 \|z_{i-1}^\parallel\| + \alpha_4 \|z_{i-1}^\perp\| \\
&\le \left(\alpha_2 + \alpha_2\alpha_4 + \alpha_2\alpha_4^2 + \cdots\right) \cdot \max_{j<i} \|z_j^\parallel\| \\
&= \frac{\alpha_2}{1-\alpha_4} \cdot \max_{j<i} \|z_j^\parallel\|.
\end{aligned}$$

The second inequality is by Lemma 3.3. The third inequality follows by induction and the fact that $z_0^\perp = 0$.

Now let us prove that $\|z_i^\parallel\| \le \alpha_6 \cdot \max_{j<i} \|z_j^\parallel\|$, where $\alpha_6 = \alpha_1 + \frac{\alpha_2 \cdot \alpha_3}{1-\alpha_4}$.

$$\begin{aligned}
\|z_i^\parallel\| &= \left\|(E\widetilde{P} z_{i-1})^\parallel\right\| \\
&\le \left\|(E\widetilde{P} z_{i-1}^\parallel)^\parallel\right\| + \left\|(E\widetilde{P} z_{i-1}^\perp)^\parallel\right\| \\
&\le \alpha_1 \|z_{i-1}^\parallel\| + \alpha_3 \|z_{i-1}^\perp\| \\
&\le \alpha_1 \|z_{i-1}^\parallel\| + \alpha_3 \cdot \alpha_5 \max_{j<i-1} \|z_j^\parallel\| \\
&\le \left(\alpha_1 + \frac{\alpha_2 \cdot \alpha_3}{1-\alpha_4}\right) \cdot \max_{j<i} \|z_j^\parallel\|.
\end{aligned}$$

Again the second inequality is by Lemma 3.3. This way we get that

$$\|z_k^\parallel\| \le \left(\alpha_1 + \frac{\alpha_2 \cdot \alpha_3}{1-\alpha_4}\right)^k \|z_0^\parallel\| = \sqrt{d} \cdot \left(\alpha_1 + \frac{\alpha_2 \cdot \alpha_3}{1-\alpha_4}\right)^k$$

and hence

$$\left\langle z_0, \left(E\widetilde{P}\right)^k z_0 \right\rangle \le d \cdot \left(\alpha_1 + \frac{\alpha_2 \cdot \alpha_3}{1-\alpha_4}\right)^k$$

Using elementary calculus, one can prove the following facts:

1. $t(2 + \log(1/t)) \le 1$ for $0 < t \le 1/4$.
2. $\exp(x) \le 1 + x + x^2 \le 1 + 2x$ for $x \in [0, 1]$.

This implies that

$$\begin{aligned}
\alpha_1 &= \exp(t\ell) - t\ell \le 1 + t^2\ell^2 = 1 + t^2(2 + \log(1/t))^2 \le 1 + 9t^2 \log^2(1/t) \\
\alpha_2\alpha_3 &= \lambda \cdot (\exp(t\ell) - 1)^2 \le 4\lambda t^2 \ell^2 \le 36\lambda t^2 \log^2(1/t)
\end{aligned}$$

Thus

$$\left(\alpha_1 + \frac{\alpha_2 \cdot \alpha_3}{1-\alpha_4}\right)^k \leq \left(1 + 9t^2\log^2(1/t) + \frac{36\lambda}{1-\sqrt{\lambda}}t^2\log^2(1/t)\right)^k$$
$$\leq \left(1 + 9t^2\log^2(1/t) + \frac{72\lambda}{1-\lambda}t^2\log^2(1/t)\right)^k$$
$$\leq \exp\left(k\left(9t^2\log^2(1/t) + \frac{72\lambda}{1-\lambda}t^2\log^2(1/t)\right)\right)$$

This completes the proof. □

Remark 3.4. As is the case with Healy's proof [Hea08], our proof also works for the case when there are different mean zero functions $f_1, \ldots, f_k$ for the different steps of the walk and also when there are k λ-expanders $G_1, \ldots, G_k$ and the j^{th} step of the walk is taken according to G_j.

Remark 3.5. We also suspect that with appropriate modifications, our proof should generalize to random walks on irregular undirected graphs (or reversible Markov chains) as was done for Healy's proof in [CLLM12].

4 Proof of Theorem 1.4

Proof. Observe that for every $i = 2, \ldots, k$ we have

$$\mathbb{E}(f(v_i)|v_{i-1}) = \sum_{u \sim v_{i-1}} P(v_{i-1}, u)f(u) = Pf(v_{i-1}),$$

whence the random vectors

$$Y_i^{(1)} := f(v_i) - Pf(v_{i-1})$$

satisfy

$$\mathbb{E}(Y_i^{(1)}|v_1, \ldots, v_{i-1}) = 0$$

and thus form a martingale difference sequence with respect to the filtration generated by initial segments of $v_1, \ldots, v_k$. Denoting $Y_1^{(1)} := f(v_1)$, we can write the sum of interest as a martingale part plus a remainder, which is a sum of $k-1$ (i.e., one fewer) random variables:

$$S = \sum_{i=1}^{k} Y_i^{(1)} + \sum_{i=1}^{k-1} Pf(v_i).$$

Notice that Pf is also a mean zero function on G, and by Jensen's inequality we have

$$\|(Pf)(v)\|_* \leq M := \max_{v \in V} \|f(v)\|_* \qquad \text{for all } v.$$

The key point is that the remainder terms $Pf(v_i)$ are smaller on average than the original terms $f(v_i)$ in squared Euclidean norm, because P is a contraction orthogonal to the constant vector; in particular, by considering the action of P on each coordinate of f separately, we have:

$$\sum_v \|Pf(v)\|_2^2 \leq \lambda \cdot \sum_v \|f(v)\|_2^2. \tag{11}$$

Iterating this construction on the remainder a total of $T \leq k$ times, we obtain a sequence of martingales $1 \leq t \leq T$:

$$Y_1^{(t)} := P^{t-1} f(v_1) \qquad Y_i^{(t)} := P^{t-1} f(v_i) - P^t f(v_{i-1}), \qquad i = 1 \ldots, k - (t-1)$$

which are related to the original sum as:

$$S = \sum_{t=1}^{T} \sum_{i=1}^{k-t+1} Y_i^{(t)} + \sum_{i=1}^{k-T} P^T f(v_i).$$

Interchanging the order of summation, we find that the random matrices

$$Z_i := \sum_{t=1}^{\min\{(k+1-i),T\}} Y_i^{(t)}$$

themselves form a martingale difference sequence, with each

$$\|Z_i\|_* \leq \sum_t \|Y_i^{(t)}\|_* \leq TM.$$

We bound the error $W := \frac{1}{k} \sum_{i=1}^{k-T} (P^T f)(v_i)$ crudely as:

$$\|W\|_2 \leq \frac{1}{k} \sum_{i=1}^{k-T} \|P^T f(v_i)\|_2 \leq \frac{k-T}{k} \sum_{v \in V} \|(P^T f)(v)\|_2 \leq \lambda^{T/2} F \leq \exp(-(1-\lambda)T/2) F,$$

where $F := \sqrt{\sum_{v \in V} \|f(v)\|_2^2}$, by applying (11). Rearranging and setting $T = 2\log(F/\epsilon)/(1-\lambda)$ yields the advertised bound on $\|Z_i\|_*$. □

References

[AW02] Rudolf Ahlswede and Andreas Winter. Strong converse for identification via quantum channels. *IEEE Transactions on Information Theory*, 48(3):569–579, 2002.

[Che52] Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, pages 493–507, 1952.

[CL06] Fan Chung and Linyuan Lu. Concentration inequalities and martingale inequalities: a survey. *Internet Mathematics*, 3(1):79–127, 2006.

[CLLM12] Kai-min Chung, Henry Lam, Zhenming Liu, and Michael Mitzenmacher. Chernoff-Hoeffding bounds for Markov chains: Generalized and simplified. *arXiv preprint arXiv:1201.0559*, 2012.

[Gil98] David Gillman. A chernoff bound for random walks on expander graphs. *SIAM Journal on Computing*, 27(4):1203–1220, 1998.

[Gol65] Sidney Golden. Lower bounds for the Helmholtz function. *Physical Review*, 137(4B):B1127, 1965.

[Hea08] Alexander D Healy. Randomness-efficient sampling within NC1. *Computational Complexity*, 17(1):3–37, 2008.

[Hir52] Isidore Isaac Hirschman. A convexity theorem for certain groups of transformations. *Journal d'Analyse Mathématique*, 2(2):209–218, 1952.

[HKT16] Fumio Hiai, Robert Koenig, and Marco Tomamichel. Generalized log-majorization and multivariate trace inequalities. *arXiv preprint arXiv:1609.01999*, September 2016.

[Kah97] Nabil Kahale. Large deviation bounds for Markov chains. *Combinatorics, Probability and Computing*, 6(04):465–474, 1997.

[Lez98] Pascal Lezaud. Chernoff-type bound for finite Markov chains. *Annals of Applied Probability*, pages 849–867, 1998.

[LP04] C. A. Leon and F. Perron. Optimal Hoeffding bounds for discrete reversible Markov chains. *Annals of Applied Probability*, 14(2), 2004.

[LPS88] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[LT13] Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer Science & Business Media, 2013.

[RR17] Shravas Rao and Oded Regev. A sharp tail bound for the expander random sampler. *arXiv preprint arXiv:1703.10205*, 2017.

[Rud99] Mark Rudelson. Random vectors in the isotropic position. *Journal of Functional Analysis*, 164(1):60–72, 1999.

[RVW00] Omer Reingold, Salil Vadhan, and Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 3–13. IEEE, 2000.

[SBT17] David Sutter, Mario Berta, and Marco Tomamichel. Multivariate trace inequalities. *Communications in Mathematical Physics*, 352(1):37–58, 2017. arXiv:1604.03023.

[Tho65] Colin J Thompson. Inequality with applications in statistical mechanics. *Journal of Mathematical Physics*, 6(11):1812–1813, 1965.

[Tro12] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.

[Tro15] Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.

[Wag08] R. Wagner. Tail estimates for sums of variables sampled by a random walk. *Combinatorics, Probability and Computing*, 17(2), 2008.

[WX] Avi Wigderson and David Xiao. Retraction of Wigderson-Xiao: A randomness-efficient sampler for matrix valued functions, eccc tr05-107. 2006.

[WX05] Avi Wigderson and David Xiao. A randomness-efficient sampler for matrix-valued functions and applications. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 397–406. IEEE, 2005.

[WX08] Avi Wigderson and David Xiao. Derandomizing the AW matrix-valued Chernoff bound using pessimistic estimators and applications. *Theory of Computing*, 4(3):53–76, 2008.